arXiv:1603.05180v1 [physics.comp-ph] 16 Mar 2016

Linear embedding of free energy minimization

Jonathan E. Moussa*
Center for Computing Research, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA

Exact free energy minimization is a convex optimization problem that is usually approximated with stochastic sampling methods. Deterministic approximations have been less successful because many desirable properties have been difficult to attain. Such properties include the preservation of convexity, lower bounds on free energy, and applicability to systems without subsystem structure. We satisfy all of these properties by embedding free energy minimization into a linear program over energy-resolved expectation values. Numerical results on small systems are encouraging, but a lack of size consistency necessitates further development for large systems.

PACS numbers: 05.10.-a

Deterministic approximations of free energy minimization trace back to the early work of Bethe [1]. His approach was later generalized into the cluster variation method (CVM) [2] and has served as the foundation for a broad class of modern statistical inference methods [3]. Such deterministic methods have not been as successful as stochastic methods such as the Metropolis algorithm [4] because of a few persistent technical problems. Relaxation of the convex free energy minimization problem should produce a lower bound on free energy that is the unique global minimum of a simpler convex optimization problem. Only recently have methods begun to achieve this goal by restricting statistical correlations to tree [5] or causal [6] structures. A more persistent problem is that systems must be decomposable into subsystems with a moderate number of microstates. This is a problem for rotors in classical systems or bosons in quantum systems where all subsystems have an infinite number of microstates without artificial truncations.

At zero temperature, the minimization of free energy over microstate probability distributions reduces to a search for a microstate with the lowest energy. If there is a large number of microstates, then this is still a hard optimization problem. However, relaxation methods in discrete optimization such as the cutting-plane method [7] can reduce hard optimizations to computationally tractable convex optimizations. The relaxed problems optimize over supersets of the original domain and produce lower bounds on minima. While there are no known universal methods for relaxing all hard problems, a variety of problem-specific relaxation methods exist [8].

In this paper, we study a convex relaxation of free energy minimization to linear programs. The primary difficulty with applying existing relaxation techniques to this problem is the nonlinearity of entropy on microstate probabilities. All known relaxations depend on linear information inequalities [9] and a decomposition of systems into small subsystems. We derive an alternate approach that embeds each expectation value in a distribution of contributions over microstate energy. This is a generalization of the energy density of states sampled by the Wang-Landau algorithm [10]. The embedding linearizes the free energy, which enables a cutting-plane relaxation.

* godotalgorithm@gmail.com

Our goal is to compute the minimum free energy $F(T)$ at a temperature T of a system with N microstates,

$$F(T) = \min_{p_i} \sum_{i=1}^{N} (E_i p_i + T p_i \ln p_i)$$
$$\text{subject to } \sum_{i=1}^{N} p_i = 1 \text{ and } p_i \geq 0. \tag{1}$$

Microstate i has energy E_i and probability p_i. T is in units of energy. The minimizer $p_i(T)$ of Eq. (1) is the Boltzmann distribution with partition function $Z(T)$,

$$p_i(T) = \frac{\exp(-E_i/T)}{Z(T)},\ Z(T) = \sum_{i=1}^{N} \exp(-E_i/T). \tag{2}$$

$F(T)$ is not efficient to compute directly when N is large.

We assume a generic form of structure in Eq. (1) that does not make reference to subsystems. The assumption is that a system can be described by the expectation values of a small number of observables, $n \ll N$. We can then reduce p_i to n expectation values x_i and E_i to n coefficients h_i as

$$x_i = \sum_{j=1}^{N} B_{i,j} p_j \text{ and } \sum_{i=1}^{n} h_i B_{i,j} = E_j \tag{3}$$

where observables are defined by a matrix $B_{i,j}$ with linearly independent columns and $B_{1,i} = 1$. A standard example is a description of spin systems by few-spin expectation values.

To simplify Eq. (1), we must enforce $p_i \geq 0$ with only x_i. The feasible set of x_i values is hard to compute. Instead, we relax the N constraints on p_i to $m \ll N$ constraints on x_i as

$$p_i \geq 0 \text{ and } \sum_{j=1}^{n} A_{i,j} B_{j,k} \geq 0 \ \Rightarrow\ \sum_{j=1}^{n} A_{i,j} x_j \geq 0, \tag{4}$$

for an m-by-n matrix $A_{i,j}$ that identifies linear combinations of observables with guaranteed nonnegative expectation values. Eq. (4) is the central approximation of our $F(T)$ lower bound construction. A standard example is nonnegativity of marginal probabilities for few-spin subsystems of a spin system.

In the $T = 0$ limit, $F(T)$ reduces to finding the microstate i that has the lowest energy E_i. This is a linear program in N dimensions with variables p_i. We relax Eq. (1) with Eq. (4) to a linear program in n dimensions with variables x_i,

$$\begin{aligned} F(0) &\geq \min_{x_i} \sum_{i=1}^{n} h_i x_i \\ &\text{subject to } x_1 = 1 \text{ and } \sum_{j=1}^{n} A_{i,j} x_j \geq 0. \end{aligned} \tag{5}$$

In this context, Eq. (4) is a form of the cutting-plane method [7] in discrete optimization that simplifies the feasible set.

For $T > 0$, the convex relaxation of $F(T)$ in Eq. (1) is more difficult because of the nonlinearity of free energy in p_i, which is hard to approximate with x_i. Inspired by the Wang-Landau algorithm [10], we decompose p_i into an integral over energy density of states. We write this using normalized distributions $P_i(s)$ over the surprisal, $s = \ln Z(T) + E/T$, such that

$$\int_0^\infty P_i(s)\, ds = 1 \text{ and } \int_0^\infty P_i(s)\, e^{-s} ds = p_i. \tag{6}$$

The N distributions $P_i(s)$ solve a continuous linear program

$$\begin{aligned} F(T) &= \min_{P_i(s)} \sum_{i=1}^{N} \int_0^\infty (E_i - Ts) P_i(s)\, e^{-s} ds \\ &\text{subject to } \sum_{i=1}^{N} \int_0^\infty P_i(s)\, e^{-s} ds = 1,\ P_i(s) \geq 0, \\ &\text{and } \int_0^\infty P_i(s)\, ds = 1. \end{aligned} \tag{7}$$

This is a straightforward linear embedding of Eq. (1). Every normalized nonnegative distribution collapses and results in a Dirac delta function, $P_i(s) = \delta(s - \ln Z(T) - E_i/T)$.

The $T = 0$ relaxation method can now be applied at $T > 0$. First, we relax the distributional nonnegativity of $P_i(s)$,

$$P_i(s) \geq 0 \ \Rightarrow\ \sum_{j=1}^{n} A_{i,j} X_j(s) \geq 0. \tag{8}$$

Distributions $X_i(s)$ subsume x_i and Eq. (8) subsumes Eq. (4). Next, we relax Eq. (7) to the continuous linear program

$$\begin{aligned} F(T) &\geq \min_{X_i(s)} \sum_{i=1}^{n} \int_0^\infty (h_i - \delta_{1,i} Ts) X_i(s)\, e^{-s} ds \\ &\text{subject to } \int_0^\infty X_1(s)\, e^{-s} ds = 1,\ \sum_{j=1}^{n} A_{i,j} X_j(s) \geq 0, \\ &\text{and } \int_0^\infty X_i(s)\, ds = g_i := \sum_{j=1}^{N} B_{i,j}. \end{aligned} \tag{9}$$

The number of optimization variables has been reduced from N to n distributions. This bound is equivalent to Eq. (5) in the $T = 0$ limit. Further relaxation requires discretization.

We finally relax Eq. (7) to a finite linear program by testing nonnegativity against r nonnegative test functions $t_i(s)$,

$$x_{i,j} = \int_0^\infty X_i(s) t_j(s)\, ds, \tag{10}$$

that fit all integrands over s with coefficients u_i, v_i, and w_i,

$$\sum_{i=1}^{r} \begin{bmatrix} u_i \\ v_i \\ w_i \end{bmatrix} t_i(s) = \begin{bmatrix} 1 \\ e^{-s} \\ se^{-s} \end{bmatrix}, \tag{11}$$

The relaxed finite linear program with nr variables $x_{i,j}$ is

$$\begin{aligned} F(T) &\geq \min_{x_{i,j}} \sum_{i=1}^{n} h_i x_i - T \sum_{i=1}^{r} w_i x_{1,i} \\ &\text{subject to } x_1 = 1,\ \sum_{j=1}^{n} A_{i,j} x_{j,k} \geq 0, \\ &\text{and } \sum_{j=1}^{r} \begin{bmatrix} u_j \\ v_j \end{bmatrix} x_{i,j} = \begin{bmatrix} g_i \\ x_i \end{bmatrix}. \end{aligned} \tag{12}$$

This naturally reduces to Eq. (5) at $T = 0$. The cost of $T > 0$ is an increase in the number of variables from n to nr.

For the example considered in this paper, we construct test functions as piecewise linear combinations of 1, e^{-s}, and se^{-s} specifically for odd r. They are defined by a set of $(r+1)/2$ interpolation points s_i and split into two types of functions,

$$\begin{aligned} t_{2i-1}(p) &= \begin{cases} \frac{(1-s_{i-1}+s)\exp(s_{i-1}-s)-1}{(1-s_{i-1}+s_i)\exp(s_{i-1}-s_i)-1}, & s_{i-1} \leq s \leq s_i \\ \frac{(1-s_{i+1}+s)\exp(s_{i+1}-s)-1}{(1-s_{i+1}+s_i)\exp(s_{i+1}-s_i)-1}, & s_i \leq s \leq s_{i+1} \\ 0, & \text{otherwise} \end{cases}, \\ t_{2i}(s) &= \begin{cases} f_i(s)/f_i(\bar{s}_i), & s_i \leq s \leq s_{i+1} \\ 0, & \text{otherwise} \end{cases}, \\ f_i(s) &= \tfrac{s_{i+1}-s}{s_{i+1}-s_i} e^{s_i - s} + \tfrac{s-s_i}{s_{i+1}-s_i} e^{s_{i+1}-s} - 1, \\ \bar{s}_i &= 1 + \tfrac{s_i - s_{i+1}\exp(s_i - s_{i+1})}{1-\exp(s_i - s_{i+1})}, \end{aligned} \tag{13}$$

with $s_i < \bar{s}_i < s_{i+1}$. An arbitrary function $C(s)$ is interpolated by a linear combination of these functions with coefficients

$$\begin{aligned} c_{2i-1} &= C(s_i), \\ c_{2i} &= C(\bar{s}_i) - c_{2i-1} t_{2i-1}(\bar{s}_i) - c_{2i+1} t_{2i+1}(\bar{s}_i). \end{aligned} \tag{14}$$

This interpolation is exact for 1, e^{-s}, and se^{-s}, and it defines the coefficients in Eq. (11). We use uniform grids of s_i values that cover the range of possible s values in each calculation.

To demonstrate the dimensional reduction capability of our linear embedding of free energy minimization, we consider a clock with N microstates labeled by an angle $\theta_i = 2\pi(i-1)/N$. Convenient observables of the clock are n moments,

$$B_{2i+1,j} = \cos i\theta_j \text{ and } B_{2i,j} = \sin i\theta_j. \tag{15}$$

To constraint their expectation values, we must identify linear combinations of moments that are nonnegative for every θ_i. A clock with moment observables and constraints is an example of non-subsystem structure described by Eqs. (3) and (4).

We consider a simple clock Hamiltonian with two nonzero terms, $h_2 = -\sin\phi$ and $h_3 = -\cos\phi$. For ample dimensional reduction opportunities, we consider $N = 2^{20}$. We solve all of the linear programs with ECOS [11]. Eq. (5) is implemented as shown, but Eq. (12) is rescaled to improve numerical stability [12]. We use $v_j x_{i,j}$ as variables instead of $x_{i,j}$, and equations with $g_i = N\delta_{1,i}$ are rescaled by v_r. This avoids large quantities in the program and its solution, but persistent solver stability problems still limit our numerical experiments.

At $T = 0$, it is possible to achieve a tight bound with Eq. (5) using only observables that appear in the Hamiltonian with a sufficient set of constraints. For our example with x_1, x_2, and x_3, these are N constraints defined by

$$A_{i,1} = 1,\ A_{i,2} = -\sin\theta_i,\ \text{and}\ A_{i,3} = -\cos\theta_i. \tag{16}$$

With $N \gg m$, we approximate these constraints by uniformly sampling m points from θ_i. Error versus m is shown in Fig. 1 with the Hamiltonian tuned to $\phi = \pi/m$ to maximize the error. There is an m^{-2} scaling of error, which is comparable to the error in estimating the minimum of a smooth function from its evaluation at a uniform set of points without interpolation.

AT $T > 0$, we set s_1 and s_r to be the exact minimum and maximum surprisal values and achieve apparent convergence with $r = 41$. With 3 observables, the error saturates in Fig. 1. We extend these calculations to 5 and 7 observables with the addition of constraints defined by nonnegative functions of θ_i,

$$\prod_{j=1}^{q}[1 - \cos(\theta_i - \theta_{\Lambda(j)})] \geq 0, \tag{17}$$

defined by q angles, $\theta_{\Lambda(j)}$, for $q \leq (n-1)/2$. We again limit the number of constraints by uniformly sampling these angles from θ_i. A systematic reduction of error with increasing n is evident in Fig. 1 at $T > 0$. Unlike $T = 0$, we need to consider observables that do not appear in the Hamiltonian to improve the accuracy at $T > 0$. The cause of these differences between $T = 0$ and $T > 0$ is unknown and left for future studies.

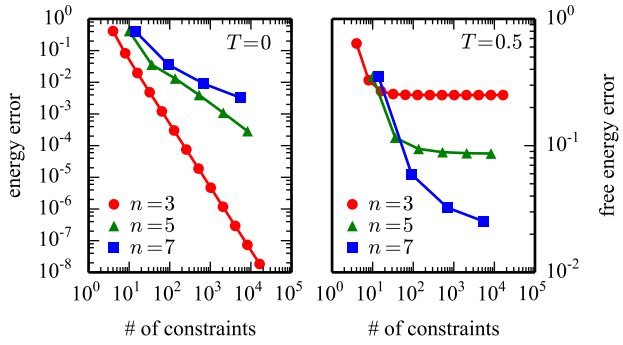

FIG. 1. (color online) Relaxed energy minimization (left) and free energy minimization (right) for a clock with $N = 2^{20}$ and microstate energies $E(\theta) = -\cos(\theta - \phi)$. Its exact energies are $F(0) = -1$ and $F(0.5) \approx -7.343$. At $T = 0$, the error converges with only $n = 3$ observables, and extra observables simply reduce the approximation efficiency. At $T = 0.5$, the error stagnates with $n = 3$ observables, and extra observables reduce the minimum attainable error.

For a many-body method to be accurate for large systems, size-consistency errors must be small [13]. These errors occur when separate and combined calculations of two independent subsystems disagree. We combine Eq. (9) for two subsystems 'L' and 'R' by merging the linear programs in a block matrix notation to partition observables ('1' separated from '> 1'),

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1^L & \mathbf{A}_{>1}^L & 0 \\ \mathbf{A}_1^R & 0 & \mathbf{A}_{>1}^R \end{bmatrix},\ \mathbf{B} = \begin{bmatrix} \mathbf{B}_1^L \otimes \mathbf{B}_1^R \\ \mathbf{B}_{>1}^L \otimes \mathbf{B}_1^R \\ \mathbf{B}_1^L \otimes \mathbf{B}_{>1}^R \end{bmatrix},$$
$$\mathbf{h} = \begin{bmatrix} h_1^L + h_1^R \\ \mathbf{h}_{>1}^L \\ \mathbf{h}_{>1}^R \end{bmatrix},\ \mathbf{X}(E) = \begin{bmatrix} X_1^L * X_1^R \\ \mathbf{X}_{>1}^L * X_1^R \\ X_1^L * \mathbf{X}_{>1}^R \end{bmatrix}(s),$$
$$(f * g)(s) = \int_0^s f(s - s')g(s')\,ds'. \tag{18}$$

By construction, the combined solution is feasible and gives a free energy bound equal to the sum of the subsystem free energy bounds. However, it is not necessarily a stable solution to Eq. (9) and any further reductions of the free energy bound after reoptimization constitute a size-consistency error.

We study the size-consistency errors of a linear embedding for independent spins in a uniform field. We model each spin exactly with two observables and two constraints,

$$2\mathbf{A} = \mathbf{B} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix},\ \mathbf{h} = \begin{bmatrix} 0 \\ h_2 \end{bmatrix}. \tag{19}$$

This system is invariant to permutations of spins, and we can reduce Eq. (9) for k spins to a single $X_1(s)$ and $X_2(s)$,

$$F(T) \geq \min_{X_2(s)} \int_0^\infty [kh_2 X_2(s) - Ts|X_2(s)|]\ e^{-s} ds$$
$$\text{subject to} \int_0^\infty \begin{bmatrix} |X_2(s)| \\ |X_2(s)|e^{-s} \\ X_2(s) \end{bmatrix} ds = \begin{bmatrix} 2^k \\ 1 \\ 0 \end{bmatrix}, \tag{20}$$

with $X_1(s) = |X_2(s)|$ set by saturating the two inequalities. The optimal $X_1(s)$ reduces to a sum of two Dirac delta functions,

$$X_1(s) = \sum_{i=0}^{1} 2^{k-1}\delta(s + F(T)/T + (2i-1)kh_2/T),$$
$$F(T) = -kT(\ln 2 + k^{-1}\ln\cosh kh_2/T). \tag{21}$$

By contrast, the exact $X_1(s)$ has $k+1$ Dirac delta functions and the exact $F(T)$ is a strictly linear function of k,

$$X_1(s) = \sum_{i=0}^{k} \frac{k!}{i!(k-i)!}\delta(s + F(T)/T + (2i-k)h_2/T),$$
$$F(T) = -kT(\ln 2 + \ln\cosh h_2/T). \tag{22}$$

In the $k \to \infty$ limit, the approximate $F(T)/k$ asymptotes to

$$F(T)/k \to -T\ln 2 - h_2. \tag{23}$$

This is accurate in both the $T = 0$ and $T \to \infty$ limits, but it has a maximum relative error of ≈ 0.5 at $h_2/T \approx 0.8$.

In its present form, linear embedding has properties that are complementary to conventional cluster-based methods such as CVM [2]. By construction, cluster expansions of entropy have no size-consistency errors. However, they are based on exact entropy calculations of small subsystems, which assumes the efficient summation over subsystem microstates. For systems composed of subsystems with a large or even infinite number of microstates, brute force cluster expansions are intractable. Examples include clocks with large N and continuous rotors in the $N \to \infty$ limit of a clock. It may be possible to combine linear embedding techniques that reduce the size and cost of subsystem calculations with cluster expansions that maintain overall size consistency. The further development of efficient but non-size-consistent subsystem methods will complement orthogonal development of cluster-based methods such as the identification of new information inequalities [9].

For linear embedding to develop into a general-purpose free energy minimization method, size-consistency errors must be eliminated. The simplest possibility would be to identify new constraints that guarantee the stability of combined solutions in Eq. (18). In addition to the structure postulated in Eqs. (3) and (4), we might attempt to impose predetermined subsystem structure. When combining calculations of two systems, their subsystem structure must also be merged. Another plausible size-consistent extension of linear embedding is to extend the s parameter used in Eq. (6) into a separate s_i parameter for each microstate that extends the temperature dependence into a full characterization of the thermodynamic state space [14]. This simplifies the combination of independent subsystems in Eq. (18) from a convolution to a product. However, it creates a high-dimensional space of variables that must be reduced to coincide with the small number of important observables. A Fisher information matrix can be defined on this space, and its spectral nonnegativity might be a useful constraint. The large number of variables introduced by this extension would need to be compensated for by a large number of constraints.

Linear embedding of free energy minimization preserves its convexity but changes its basic structure. The exact problem has a linear objective function only at zero temperature, while the approximate problem always has a linear objective. These linear problems are optimized on a boundary of their feasible set. At finite temperature, the $p_i \ln p_i$ entropy terms in Eq. (1) drive solutions away from the boundary, which are defined by $p_i = 0$ for one or more i. This fundamentally changes how the problem responds to perturbations. Responses are determined by derivatives of the objective in the nonlinear case and by the shape of the boundary in the linear case. The consequences of this difference are unclear. Any convenience gained from the use of linear program solvers for free energy minimization is possibly offset by a fundamental distortion of the problem. It is still possible for nonlinear embedding to retain the property of having extrema only on the interior of the feasible set.

There are also computational complexity differences in free energy minimization between zero and finite temperature. At zero temperature, the search for the lowest-energy microstate is NP-hard. In general, accurate approximation of a solution is intractable on a classical computer in polynomial time [15]. At finite temperature, free energy approximation is equivalent to approximation of #P-hard counting problems [16]. While a complete understanding of the relative complexity between zero and finite temperature is not yet known, it is possible to approximate #P problems with access to an NP oracle [17].

Ultimately, this paper is meant to motivate and inspire new ideas in deterministic methods for free energy minimization. We have developed a linear embedding concept, computed a numerical example, presented an analytically solvable model, and identified the problem of size-consistency errors. While it is of limited applicability in its present form, there are many possible directions of future development. In particular, there are multiple possible ways to combine linear embedding with cluster-based methods to fix size-consistency errors. Finally, we note that linear embedding has a straightforward quantum generalization as do many cluster-based methods [6]. We need methods such as linear embedding to simulate the unbounded Hilbert spaces of bosons without artificial truncations.

I thank Norm Tubman for useful discussions. This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

[1] H. A. Bethe, Proc. R. Soc. A **150**, 552 (1935).
[2] R. Kikuchi, Phys. Rev. **81**, 988 (1951).
[3] M. J. Wainwright and M. I. Jordan, Found. Trends Machine Learn. **1**, 1 (2008).
[4] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
[5] M. Wainwright, T. S. Jaakkola, and A. Willsky, IEEE Trans. Inf. Theory **51**, 2313 (2005).
[6] D. Poulin and M. B. Hastings, Phys. Rev. Lett. **106**, 080403 (2011).
[7] J. E. Kelley, Jr., J. Soc. Ind. Appl. Math. **8**, 703 (1960).
[8] M. Laurent, Math. Oper. Res. **28**, 470 (2003).
[9] R. W. Yeung, IEEE Trans. Inf. Theory **43**, 1924 (1997).
[10] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[11] A. Domahidi, E. Chu, and S. Boyd, in *2013 European Control Conference (ECC)*, pp. 3071-3076 (IEEE, New York, 2013).
[12] See ancillary files at *pending approval* for source code.
[13] J. A. Pople, J. S. Binkley, and R. Seeger, Int. J. Quantum Chem., Symp. **10**, 1 (1976).
[14] G. E. Crooks, Phys. Rev. Lett. **99**, 100602 (2007).
[15] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy, J. ACM **45**, 501 (1998).
[16] M. Singh and N. K. Vishnoi, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pp. 50-59 (ACM, New York, 2014).
[17] L. Stockmeyer, SIAM J. Comput. **14**, 849 (1985).